


NLIT 2007: Fermilab Site Update

Mark O. Kaletka

Associate Division Head for Laboratory & Scientific Core Services

*Fermi National Accelerator Laboratory
Office of Science / U.S. Department of Energy
Managed by Fermi Research Alliance LLC*

The Quantum Universe



- *Einstein's Dream of Unified Forces*
 - 1. Are there undiscovered principles of nature: new symmetries, new physical laws?
 - 2. How can we solve the mystery of dark energy?
 - 3. Are there extra dimensions of space?
 - 4. Do all the forces become one?
- *The Particle World*
 - 5. Why are there so many kinds of particles?
 - 6. What is dark matter? How can we make it in the laboratory?
 - 7. What are neutrinos telling us?
- *The Birth of the Universe*
 - 8. How did the universe come to be?
 - 9. What happened to the antimatter?

Science @Fermilab




- Fermilab has four main efforts to help answer these questions:
 - The Tevatron collider experiments (CDF, DZero)
 - Experiments in neutrino physics (MINOS, MiniBOONE)
 - Experimental astrophysics (SDSS, Auger, CDMS)
 - The CMS experiment at the CERN LHC (USCMS)
- As well as preparations for new accelerators and detectors (ILC, DES, NOvA)


A Time of Change



- Rampdown of the Tevatron program
- Preparations for data taking by US CMS at the Large Hadron Collider (LHC) at CERN
 - Fermilab is the US CMS Tier 1 regional computing center
- Participation in the design and planning for the International Linear Collider
 - 2 x 40km electron linear accelerators



Our Mission




- Accomplished by huge worldwide collaborations
- Produces petabytes of data
- Requires enormous computing resources that can only be met by global grids
- Interchanges data, ideas, people worldwide on a daily basis
- Leads other sciences in making grid computing as important and useful as the Web (invented in HEP)

Scientific Computing @ Fermilab



- Data flows from the experiments into mass storage
 - ~ 6 PB in tape storage
- Bulk production processing & physics analysis use grid-enabled compute farms
 - ~ 6000 CPU's @Fermilab = ~ 16 THz
- All tied together by high-performance networks
 - Peak off-site transfer rates ~ 5 GB/s

Open Science Grid




- FNAL is a major contributor in world-wide Grid computing and a consortium member and leader in Open Science Grid
 - Open Science Grid is a distributed computing infrastructure for large-scale scientific research, built and operated by a consortium of universities, national laboratories, scientific collaborations and software developers.
- High-Energy Physics has adopted the Grid as a major strategy for providing the massive computing required for the LHC and beyond.

FermiGrid



- Fermilab has undertaken the strategy of placing all of its production resources in a Grid "meta-facility" infrastructure called FermiGrid.
- This strategy is designed to optimize use of resources and allow Fermilab to:
 - Insure that the large experiments have first priority usage of dedicated resources purchased on their behalf.
 - Allow opportunistic use of dedicated and shared resources by Virtual Organizations (VO's) that participate in the Fermilab experimental program and by certain VO's that use the Open Science Grid (OSG)
 - Fully support Open Science Grid and the LHC Computing Grid and gain positive benefit from this infrastructure in the US and Europe

Mass Storage & Data Movement

- Enstore is used as the tape backend for storage of scientific data
 - Presents a file-system view of tape storage
 - Routinely move >30 TB per day in and out of Enstore
- dCache is used as a high-performance disk cache for transient data
 - May be used w/ or w/o Enstore
 - Provides Grid interfaces
 - Supports many replicas for performance or reliability
 - Build from commodity disk arrays (SATABeast)
- Both are joint projects involving High-Energy Physics and Grid collaborators



Networking




- Large-scale data movement requires good networks – both on- and off-site
 - We rely on ESNet for commodity & production traffic & general Internet access
 - High-impact data transfers for CMS use a dedicated path via Starlight and USLHCnet to CERN
 - Charts to left show ability to achieve 1PB/mo transfer via ESNet & Starlight during CMS “data challenges”

Facilities

■ Facilities and infrastructure are a challenge

- High-density Grid computing places special demands on power, cooling
- Most computing now located in converted experimental halls – where ample power is available
- Expanding these on a “just in time” basis (see chart)
- Main computing center – purpose-built for mainframes – is not adequate for Grids and suffers aging infrastructure (20+ yrs old)
- Enforce a rigorous 3 yr retirement cycle

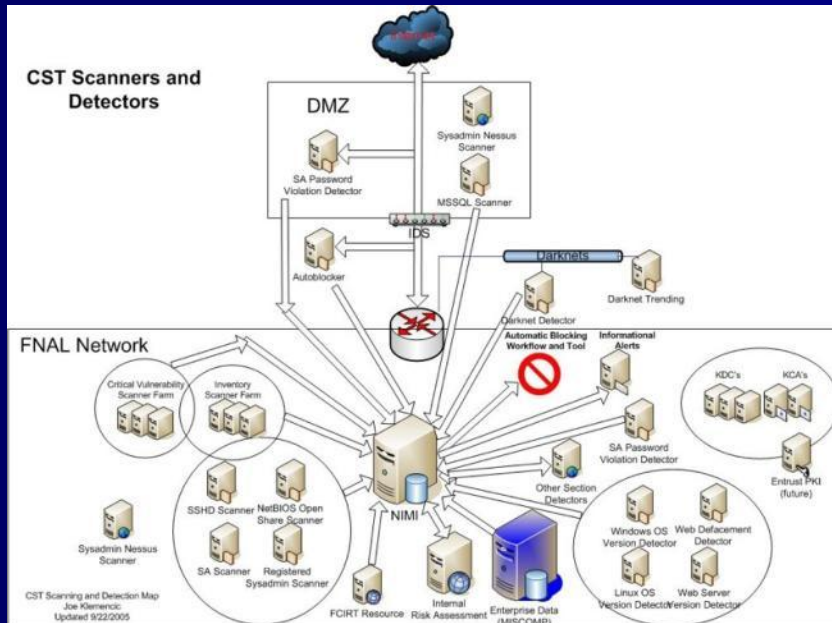


LHC@FNAL



- The CERN LHC accelerator is nearing commissioning; the locus of activity in High Energy Physics will move there
 - The Tevatron at FNAL will turn off in about 2009
- LHC@FNAL is a facility for remote monitoring & operation of the LHC accelerator & CMS experiment
 - Allows US physicists to be fully engaged in LHC & CMS operations w/o traveling to Europe
 - Emphasis on telepresence, including high-quality video conferencing
 - High visibility effort here & at CERN

Computer Security



- Challenge is to maintain a balance of security and openness in support of open science
- Risk-based program follows NIST standards
- An array of scanners and detectors with a central database (NIMI)
 - Tracks every system connected to the FNAL network
 - Identifies the sysadmin of every system
 - Scans continuously & periodically for services and vulnerabilities
 - Detects network anomalies
 - Notifies and blocks non-compliant systems
- Central Laboratory-wide authentication system
 - Kerberos- & Windows-based
 - Kerberos-derived X.509 certificates

Central Services


- Full gamut of IT services:
 - Email
 - Including spam & virus filtering
 - Sun gateways & IMAP servers
 - Web
 - Apache & IIS servers
 - Plone content & workflow management
 - Calendaring & Meeting Collaboration
 - MeetingMaker, CERN InDiCo, SPMS
 - Storage & Backup
 - BlueArc SAN/NAS, TiBS backup, OpenAFS global file system
 - HelpDesk

Desktop Support



- Windows support uses the “usual” array of Microsoft tools:
 - Single Windows domain (tied to central authentication)
 - SMS for deployment, configuration management, inventory
 - Symantec antivirus
- Mac OS support is nascent
 - Active community self-supports
 - Evaluating SMS-like tools for Mac OS

Scientific Linux



- Scientific Linux (SL) is a joint project between Fermilab, CERN and other contributors which provides an open-source distribution of Linux for the scientific (primarily High-Energy Physics) community
- Scientific Linux Fermi (SLF) provides Fermilab-specific customizations
- SL and SLF are community-supported (primarily via mailing lists)
- SLF provides infrastructure for patching, inventory and configuration management
- Some applications (primarily Oracle) require commercially-supported Red Hat Linux
- See: <https://www.scientificlinux.org/>

Summary

- Fermilab delivers a computing program which provides high-performance computing and networking as well as traditional IT services for a very demanding physics program.
- This computing program is well prepared for the transition to the CERN LHC era and beyond.

The End

